



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Adress: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

| | | | | |
|--|---------------|----------------------|---------------------|------------------|
| APPLICATION NO. | FILING DATE | FIRST NAMED INVENTOR | ATTORNEY DOCKET NO. | CONFIRMATION NO. |
| 10/749,752 | 12/30/2003 | Matthew Mattina | 42P17893 | 9070 |
| 59796 | 7590 | 09/08/2008 | EXAMINER | |
| INTEL CORPORATION c/o INTELLEVATE, LLC P.O. BOX 52050 MINNEAPOLIS, MN 55402 | | | ELAND, SHAWN | |
| ART UNIT | PAPER NUMBER | | 2188 | |
| MAIL DATE | DELIVERY MODE | | | |
| 09/08/2008 | PAPER | | | |

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

| | | |
|------------------------------|--------------------------------------|---------------------------------------|
| Office Action Summary | Application No. 10/749,752 | Applicant(s) MATTINA ET AL. |
| | Examiner SHAWN ELAND | Art Unit 2188 |

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
 - If no period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
 - Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133).
- Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) Responsive to communication(s) filed on 18 June 2008.
- 2a) This action is FINAL. 2b) This action is non-final.
- 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) Claim(s) 1 and 3-20 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) Claim(s) _____ is/are allowed.
- 6) Claim(s) 1 and 3-20 is/are rejected.
- 7) Claim(s) _____ is/are objected to.
- 8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) The specification is objected to by the Examiner.
- 10) The drawing(s) filed on _____ is/are: a) accepted or b) objected to by the Examiner.
 Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
 Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) Notice of References Cited (PTO-892)
 2) Notice of Draftsperson's Patent Drawing Review (PTO-948)
 3) Information Disclosure Statement(s) (PTO/DS/06)
 Paper No(s)/Mail Date _____
- 4) Interview Summary (PTO-413)
 Paper No(s)/Mail Date _____
- 5) Notice of Informal Patent Application
 6) Other: _____

DETAILED ACTION

A request for continued examination under 37 CFR 1.114, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.114, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 37 CFR 1.114. Applicant's submission filed on 06/18/08 has been entered.

Status of Claims

Claims 1 & 3 – 20 are pending in the Application.

Claims 1, 3 – 11, 13 – 15, & 18 – 20 have been amended.

Claim 2 is cancelled.

Claims 1 & 3 – 20 are rejected.

Response to Arguments

Applicant's arguments filed 06/18/08 have been fully considered but they are not persuasive.

Applicant argues neither Bordaz nor Jennings disclose a single message invalidating a block in one processor core and providing a write acknowledgement to another core. Applicant also argues that neither Bordaz nor Jennings disclose a single message invalidating all cores other than a requesting core.

In response to applicant's argument that the references fail to show certain features of applicant's invention, it is noted that the features upon which applicant relies (i.e., a single message invalidating all cores other than a requesting core) are not recited in the rejected claim(s). Although the claims are interpreted in light of the specification, limitations from the specification are not read into the claims. See *In re Van Geuns*, 988 F.2d 1181, 26 USPQ2d 1057 (Fed. Cir. 1993).

Applicant argues neither Bordaz nor Jennings disclose generating an evict message to an owning core for a block when a read request referencing the block is received from another core.

This is beyond the scope of the claim language. Line 8 of claim 14 states that the shared memory is "to generate" an evict message. This phrase is intended use. A recitation of the intended use of the claimed invention must result in a structural difference between the claimed invention and the prior art in order to patentably distinguish the claimed invention from the prior art. If the prior art structure is capable of performing the intended use then it meets the claim.

Applicant argues Bordaz does not disclose a cache line or block being capable of being held in all of the four listed states, nor does he disclose the notion of a custodian versus an owner.

This is also beyond the scope of the claim language. Line 7 of claim 18 states that the plurality of blocks is "capable of being held" in one of the four states. This is an intended use phrase. A recitation of the intended use of the claimed invention must result in a structural

difference between the claimed invention and the prior art in order to patentably distinguish the claimed invention from the prior art. If the prior art structure is capable of performing the intended use then it meets the claim.

Claim Objections

Claim 5 is objected to because of the following informalities:

There are two periods in the claim. One is in the middle of line 5, the other period is at the end of line 6.

Appropriate correction is required.

Claim Rejections - 35 USC § 112

The following is a quotation of the first paragraph of 35 U.S.C. 112:

The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same and shall set forth the best mode contemplated by the inventor of carrying out his invention.

Claims 5 is rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the written description requirement. The claim(s) contains subject matter which was not described in the specification in such a way as to reasonably convey to one skilled in the relevant art that the inventor(s), at the time the application was filed, had possession of the claimed invention.

Line 2 of claim 5 says and InvalidateAndAcknowledge message is sent when the block is not present and there is no custodian. There is nothing in the specification that sends the

InvalidateAndAcknowledge message during the not present/no custodian state and hence, this is new matter. The InvalidateAndAcknowledge message is only sent in response to a write and the block is present, not owned, and there is a custodian, according to the specification ([0020]). Also according to the specification, a WriteAcknowledgement is sent when the block is not present ([0019]).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

- (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

Claims 1, 3 – 6, 8 – 10 are rejected under 35 U.S.C. 103(a) as being unpatentable over *Bordaz et al.* (US Patent 6,195,728 B1), and in further view of *Jennings* (US Patent 6,134,631).

In regard to claims 1 and 18 - 20, Bordaz teaches:

a plurality of processor cores (Fig. 1, elements 1-4, 21-24, 41-44 and 61-64 depict a plurality of processor cores), **wherein the plurality of processor cores each include a private cache** (each processor contains its own private cache as depicted in Fig. 1 (element 11 is the private cache for processor 1 for example));

a shared cache to be shared by the plurality of processor cores (Fig. 1, memory (element 5) is shared among at least two processors) **and include a plurality of blocks, each of the plurality of blocks capable of being held in a not present state, a present**

and owned by a core of the plurality of cores state, a present, not owned, and custodian is a core of the plurality of core states, and a present, not owned, and no custodian state (the tags include information to indicate if the data is valid and if it is held exclusively by a particular processor – col. 5, lines 21-51; also, the following language is intended use: "to be coupled together" (lines 1 – 2), "to be associated" (line 4), "to be accessible" (line 5), "capable of being held" (line 7), and "to hold elements" (line 12); A recitation of the intended use of the claimed invention must result in a structural difference between the claimed invention and the prior art in order to patentably distinguish the claimed invention from the prior art. If the prior art structure is capable of performing the intended use then it meets the claim. A user could easily write software to hold the blocks in the above states);

wherein each of the plurality of blocks is home location for a subset of a physical address space (col. 4, lines 28-46 – the RC (element 15) makes up a portion of the total physical memory of memory element 5);

Note, Bordaz discloses cache memories 5, 15, 45 and 65 which comprise respective remote caches 15, 35, 55 and 77, and the remaining areas of memories 5, 15, 45 and 65 (5', 15', 45' and 65' respectively). In col. 4, line 28 though col. 5, line 7, Bordaz clearly teaches the address space of memories 5', 15', 45' and 65' as being either local or remote (with respect to each memory module 10, 20, 40 or 60). In other words, these memories 5, 15, 45 and 65 are used to store both local and remote data (i.e. shared by another processor from a different memory module). Further evidence that cache memories are "shared" by processors from other modules is seen in the abstract and col. 3, lines 18-47 (invention is aimed at improving cache coherency

of a plurality of modules. Data coherence would not be possible if cache memories were not shared among modules.

wherein the shared cache is to generate a first message to invalidate the block in a second processor core of the plurality of processor cores and provide a write acknowledgement to a requesting processor core, in response to receiving a write request referencing a block from the requesting processor core and the block not being owned (this system is capable of sending messages back and forth, even if the user has to write software to do it; see also “Response to Arguments section above);
an unbuffered bi-directional ring to connect the plurality of processor cores and the shared cache (Fig. 1, element 16 – the ring is used for communication between each module (elements 10, 20, etc), **the ring to transmit the first message to the requesting processor core and second processor core.**

Bordaz fails to specifically teach these particular elements as being stored on an integrated circuit (i.e. single processor chip). More specifically, Bordaz teaches four discrete modules (Fig. 1, elements 10, 20, 40 and 60) which each comprises multiple processors (each with a unique private cache), and a shared cache.

Jennings teaches a non-volatile memory with embedded programmable controller in which his plurality of modules may all implemented on a single integrated chip (storage system 50 (Fig. 1) may be a multi-chip module, or a single integrated circuit – col. 3, lines 52-58).

It would have been obvious to one of ordinary skill in the art at the time of the invention for Bordaz to implement his discrete modules on a single integrated circuit as taught by Jennings. By doing so, Bordaz could exploit the well-known benefits of single chip integration, which

includes lower manufacturing costs, and increased communication speed between the discrete elements implement on the one chip.

It is worthy to note that though Bordaz teaches an “on-module” cache rather than an on-chip cache as recited by Applicant, this would have been obvious over Bordaz as once all four modules are implemented on a single chip as discussed above in the combined teachings of Bordaz and Jennings. More specifically the shared caches within each module would be stored on that very single chip when Bordaz and Jennings are combined; hence they are “on-chip” cache. It is additionally worthy to note that the shared cache within each module acts as a system memory for storing element held by the shared memory.

As for claim 3, Bordaz teaches **wherein the shared cache includes one or more cache banks** (inherently all cache memory must be arranged in a configuration of at least one bank). Additionally, Bordaz indicates that each shared cache contains a remote access cache (RC – element 15), which is a separate memory bank within the shared cache (element 5), **wherein the one or more cache banks is responsible for a subset of a physical address space of the system, and wherein the block is associated with a physical address of the physical address space of the system** (col. 4, lines 28-46 – the RC (element 15) makes up a portion of the total physical memory of memory element 5).

As for claim 4, Bordaz teaches **wherein the first message includes an InvalidateAndAcknowledge message** (in order for cache coherency to work, the system would need to know when to invalidate cached copies and to let a processor know when to write; see “Response to Arguments” section above), **and wherein the shared cache is to generate the InvalidateAndAcknowledge message** (the “to generate” portion of this claim means the shared

cache only needs to be capable of this as this is intended use), **further in response to the block being present in the shared cache and the second processor core being a custodian for the block** (a custodian is merely a single processor that has a copy of the block but does not own it; see [0020] – [0021] of the specification).

As for claims 5 and 6, Bordaz teaches **wherein the first message includes an InvalidateAllAndAcknowledge message** (in order for cache coherency to work, the system would need to know when to invalidate cached copies and to let a processor know when to write; see "Response to Arguments" section above), **and wherein the shared cache is to generate the InvalidateAndAcknowledge message** (the "to generate" portion of this claim means the shared cache only needs to be capable of this as this is intended use), **further in response to the block not being present in the shared cache and none of the plurality of processor cores being a custodian for the block** (a custodian is merely a single processor that has a copy of the block but does not own it; see [0020] – [0021] of the specification), **wherein the processor cores are write-thru**, which **write data through to the shared cache** (col. 7, lines 56-65 – Bordaz discusses a write through cache mechanism which writes to reserved zones in the shared cache (i.e. element 25)).

As for claim 8, Bordaz teaches **wherein the shared cache is to fetch a second block from a memory and generate a write acknowledge message to provide a write acknowledgement to the requesting processor core in response to receiving a second write request referencing the second block, the second block not being present in the shared cache and not being owned by any of the plurality of processor cores** (if the block isn't

Art Unit: 2188

owned, and it's not present, then there's no reason not to allow the requesting processor write privileges).

As for claim 9, Bordaz teaches wherein the shared cache is to generate an evict message to evict a third block from an owning processor core and generate a second write acknowledge message to provide a second write acknowledgement to the requesting processor core in response to receiving a third write request referencing the third block, the third block being present in the shared cache and the owning processor core of the plurality of cores owns the third block (the "to generate" portion of this claim means the shared cache only needs to be capable of this as this is intended use; see "Response to Arguments" above).

As for claim 10, Bordaz teaches wherein a bank of the shared cache is to be a home location for a non-overlapping portion of a physical address space associated with the block (col. 4, lines 28-46 – the RC (element 15) makes up a portion of the total physical memory of memory element 5; this is just the level of associativity, which all caches have).

Claims 7 and 11 are rejected under 35 U.S.C. 103(a) as being unpatentable over the combined teachings of *Bordaz* (US Patent 6,195,728 B1) and *Jennings* (US Patent 6,134,631) as applied to claim 1 above, and in further view of *Fletcher* (US Patent 4,445,174).

As for claims 7 and 11, though Bordaz teaches wherein the plurality of processor cores include a buffer, he fails to teach the buffer as functioning merge buffer capable of purging stored data to a shared cache, and wherein each of the merge buffers are to coalesce multiple

stores to a same block, and wherein each private cache of the plurality of cores are not to hold dirty data, and wherein the buffers are only to hold dirty data.

Fletcher however teaches a multiprocessor system including a shared cache which a processor's private cache (Fig. 1, element 8) continuously stores data (permitting the merging of data (i.e. line by line) into the private memory from the main memory until an eviction is requested) – col. 1, line 62-68, and then moves the lines directly from a private cache to the shared cache, while circumventing the system's main memory (col. 2, lines 56-64).

Fletcher further discloses the private cache, which is used to merge data from the memory line by line, as coalescing multiple lines to a same block of the shared cache – col. 3, line 17-25 – copies of the same shared memory block may exist simultaneously in each private cache. In other words, data stored in a processor's private cache can exist as one memory block of the shared memory.

It would have been obvious to one of ordinary skill in the art at the time of the invention for the combined teachings of Bordaz and Jennings to further include Fletcher's multiprocessor system including a shared cache to his own system. By doing so, would realize improved system performance by having a means of automatically detecting lines of information moved to the shared cache, hence eliminating "pingponging" of lines between requesting processors as taught by Fletcher in col. 2, lines 49-65.

Claims 12 and 13 are rejected under 35 U.S.C. 103(a) as being unpatentable over the combined teachings of *Bordaz* (US Patent 6,195,728 B1) and *Jennings* (US Patent 6,134,631) as applied to claim 1 above, and in further view of *Koenen* (US PG Publication 2004/0019891 A1).

As for claims 12 and 13, though Bordaz teaches connecting his each processor module via a ring configuration as claimed by Applicant in claim 1, he fails to specifically teach the ring configuration as recited by Applicant in claims 12-13 of the pending Application.

Koenen however teaches an apparatus for optimizing performance in a multi-processing system, which includes connecting a plurality of module nodes via a synchronous, unbuffered, bi-directional ring with a fixed deterministic latency as recited by Applicant in claim 12-13. Referring to Fig. 1, a plurality of processing nodes (elements 12, 14 and 16) are connected for bi-directional communication (elements 12J, 14J and 16J) with the interconnect fabric (element 18). Note Koenen describes the fabric as including a ring structure in paragraph 0019, lines 9-12. The ring functions without the aid of a buffering system (i.e. unbuffered), and supports synchronous connections with a minimum static latency around the ring (paragraph 0026, lines 7-12 – the minimum latency is static). Furthermore, paragraph 0023 (and subsequently Table 1), describe preset latencies between each node depending on the number of nodes included in the system. With this table, the overall latency of the entire ring interconnect is known (likewise, fixed), which allows the system to synchronize communication between nodes.

It would have been obvious to one of ordinary skill in the art at the time of the invention, for the combined teachings of Bordaz and Jennings to implement Koenen's apparatus for optimizing performance in a multi-processing system. By doing so, they would benefit by using a superior interconnection fabric (as shown by Koenen in Fig. 1, element 18) for his processing

modules, which in turn could help Bordaz's NUMA machine by reducing access latency and increase system performance as taught by Koenen in paragraph 0011, lines 1-15.

Claims 14 and 16 – 17 are rejected under 35 U.S.C. 103(a) as being unpatentable over *Bordaz* (US Patent 6,195,728 B1) in view of *Jennings* (US Patent 6,134,631), and in further view of *Fletcher* (US Patent 4,445,174).

In regard to claim 14, Bordaz teaches:

a plurality of cores (Fig. 1, elements 1-4, 21-24, 41-44 and 61-64 depict a plurality of processor cores) **and a shared memory, to be accessible by each of the plurality of cores** (Fig. 1, memory (element 5) is shared among at least two processors), **connected in a ring** (Fig. 1, element 16 – the ring is used for communication between each module (elements 10, 20, etc),

wherein the plurality of processor cores each include a private cache (each processor contains its own private cache as depicted in Fig. 1 (element 11 is the private cache for processor 1 for example)),

and wherein the shared memory is to generate an evict message referencing an address to an owning processor core of the plurality of cores in response to receiving a read request referencing the address from a requesting core of the plurality of cores and the owning processor core owning a block associated with the address (this system is capable of sending messages back and forth, even if the user has to write software to do it; see also "Response to Arguments section above),

Note, Bordaz discloses cache memories 5, 15, 45 and 65 which comprise respective remote caches 15, 35, 55 and 77, and the remaining areas of memories 5, 15, 45 and 65 (5', 15', 45' and 65' respectively). In col. 4, line 28 though col. 5, line 7, Bordaz clearly teaches the address space of memories 5', 15', 45' and 65' as being either local or remote (with respect to each memory module 10, 20, 40 or 60). In other words, these memories 5, 15, 45 and 65 are used to store both local and remote data (i.e. shared by another processor from a different memory module). Further evidence that cache memories are "shared" by processors from other modules is seen in the abstract and col. 3, lines 18-47 (invention is aimed at improving cache coherency of a plurality of modules. Data coherence would not be possible if cache memories were not shared among modules.

Bordaz fails to specifically teach these particular elements as being stored on an integrated circuit (i.e. single processor chip). More specifically, Bordaz teaches four discrete modules (Fig. 1, elements 10, 20, 40 and 60) which each comprises multiple processors (each with a unique private cache), and a shared cache.

Jennings teaches a non-volatile memory with embedded programmable controller in which his plurality of modules may all implemented on a single integrated chip (storage system 50 (Fig. 1) may be a multi-chip module, or a single integrated circuit – col. 3, lines 52-58).

It would have been obvious to one of ordinary skill in the art at the time of the invention for Bordaz to implement his discrete modules on a single integrated circuit as taught by Jennings. By doing so, Bordaz could exploit the well-known benefits of single chip integration, which includes lower manufacturing costs, and increased communication speed between the discrete elements implement on the one chip.

It is worthy to note that though Bordaz teaches an “on-module” cache rather than an on-chip cache as recited by Applicant, this would have been obvious over Bordaz as once all four modules are implemented on a single chip as discussed above in the combined teachings of Bordaz and Jennings. More specifically the shared caches within each module would be stored on that very single chip when Bordaz and Jennings are combined; hence they are “on-chip” cache. It is additionally worthy to note that the shared cache within each module acts as a system memory for storing element held by the shared memory.

Finally, though Bordaz **teaches wherein the plurality of processor cores include a buffer**, he fails to teach the buffer as functioning **merge buffer** capable of purging stored data to a shared cache.

Fletcher however teaches a multiprocessor system including a shared cache which a processor’s private cache (Fig. 1, element 8) continuously stores data (permitting the merging of data (i.e. line by line) into the private memory from the main memory until an eviction is requested) – col. 1, line 62-68, and then moves the lines directly from a private cache to the shared cache, while circumventing the system’s main memory (col. 2, lines 56-64).

Fletcher further discloses the private cache, which is used to merge data from the memory line by line, as coalescing multiple lines to a same block of the shared cache – col. 3, line 17-25 – copies of the same shared memory block may exist simultaneously in each private cache. In other words, data stored in a processor’s private cache can exist as one memory block of the shared memory.

It would have been obvious to one of ordinary skill in the art at the time of the invention for the combined teachings of Bordaz and Jennings to further include Fletcher’s multiprocessor

system including a shared cache to his own system. By doing so, would realize improved system performance by having a means of automatically detecting lines of information moved to the shared cache, hence eliminating “pingponging” of lines between requesting processors as taught by Fletcher in col. 2, lines 49-65.

As for claim 16, the shared memory is a shared cache including a plurality of blocks, mad wherein the shared cache is capable of holding each of the plurality of blocks in a cache coherency state (tags are stored and associated with blocks of the cache to indicate which blocks are held exclusively (i.e. to maintain coherency) by a processor - col. 5, lines 21-51).

As for claim 17, wherein the cache coherency state for each of the plurality of blocks is selected from a group consisting of (1) a not present state, (2) a present and owned by a core of the plurality of cores state, (3) a present, not owned, and custodian is a core of the plurality of core states, and (4) a present, not owned, and no custodian state (the tags include information to indicate if the data is valid and if it is held exclusively by a particular processor – col. 5, lines 21-51).

Claim 15 is rejected under 35 U.S.C. 103(a) as being unpatentable over the combined teachings of *Bordaz* (US Patent 6,195,728 B1), *Jennings* (US Patent 6,134,631), and *Fletcher* (US Patent 4,445,174) as applied to claim 14 above, and in further view of *Koenen* (US PG Publication 2004/0019891 A1).

As for claim 15, though Bordaz teaches connecting his each processor module via a ring configuration as claimed by Applicant in claim 14, he fails to specifically teach the ring configuration as recited by Applicant in claim 15 of the pending Application.

Koenen however teaches an apparatus for optimizing performance in a multi-processing system, which includes connecting a plurality of module nodes via a synchronous, unbuffered, bi-directional ring with a fixed deterministic latency as recited by Applicant in claim 15.

Referring to Fig. 1, a plurality of processing nodes (elements 12, 14 and 16) are connected for bi-directional communication (elements 12J, 14J and 16J) with the interconnect fabric (element 18). Note Koenen describes the fabric as including a ring structure in paragraph 0019, lines 9-12. The ring functions without the aid of a buffering system (i.e. unbuffered), and supports synchronous connections with a minimum static latency around the ring (paragraph 0026, lines 7-12 – the minimum latency is static). Furthermore, paragraph 0023 (and subsequently Table 1), describe preset latencies between each node depending on the number of nodes included in the system. With this table, the overall latency of the entire ring interconnect is known (likewise, fixed), which allows the system to synchronize communication between nodes.

It would have been obvious to one of ordinary skill in the art at the time of the invention, for the combined teachings of Bordaz, Jennings and Fletcher to implement Koenen's apparatus for optimizing performance in a multi-processing system. By doing so, they would benefit by using a superior interconnection fabric (as shown by Koenen in Fig. 1, element 18) for his processing modules, which in turn could help Bordaz's NUMA machine by reducing access latency and increase system performance as taught by Koenen in paragraph 0011, lines 1-15.

Examiner's Information

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Shawn Eland whose telephone number is (571) 270-1029. The examiner can normally be reached on MO - TH, & every other FR.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Hyung Sough can be reached on (571) 272-4199. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Shawn Eland/
Examiner, Art Unit 2188
9/5/2008

/Hyung S. Sough/
Supervisory Patent Examiner, Art Unit 2188
09/02/08